arXiv:1701.05177v1 [stat.AP] 18 Jan 2017

Statistical Power in Longitudinal Network Studies

Christoph Stadtfeld* [1], Tom A. B. Snijders[2,3], Christian Steglich[2] and Marijtje van Duijn[2]

[1]Chair of Social Networks, ETH Zürich, Switzerland
[2]Department of Sociology, University of Groningen, Netherlands
[3]Nuffield College, University of Oxford, UK

Abstract

Longitudinal social network studies may easily suffer from a lack of statistical power. This is the case in particular for studies that simultaneously investigate change of network ties and change of nodal attributes. Such selection and influence studies have become increasingly popular due to the introduction of stochastic actor-oriented models (SAOMs). This paper presents a simulation-based procedure to evaluate statistical power of longitudinal social network studies in which SAOMs are to be applied. It describes how researchers can test different possible research designs decisions (e.g., about network delineation and study time) under uncertainty about the prevalence and strength of various social mechanisms. Two detailed case studies illustrate that network size, number of data collection waves, effect sizes, missing data, and participant turnover can have a serious effect on the statistical power of longitudinal social network studies.

*Corresponding author: Christoph Stadtfeld, ETH Zürich, Chair of Social Networks, Department of Humanities, Social and Political Sciences, Clausiusstrasse 50, 8092 Zürich, Switzerland, phone: +41 44 632 07 93, e-mail: c.stadtfeld@ethz.ch

1 Introduction

Longitudinal social network studies are costly and time-consuming both for researchers and participants. A lack of significant statistical evidence for a hypothesis should thus not originate from a research design that was "just too small" and, therefore, has insufficient statistical power (Cohen, 1977).

The introduction of Stochastic Actor-Oriented Models for the simultaneous investigation of network and behavior changes (SAOMs, Snijders et al., 2010b; Steglich et al., 2010) enabled a large number of publications that empirically study selection processes (changes in social relations in response to individual attributes) and influence processes (changes in individual attributes in response to social relations). SAOMs are typically applied to network panel data (a set of interconnected individuals surveyed in multiple data collection waves) and evaluate dynamic tendencies of individuals to change (add or drop) network ties and to change (increase or decrease) some type of behavior or individual attribute. Veenstra et al. (2013) review a number of those studies[1] and report mixed evidence regarding the prevalence of selection and influence mechanisms in adolescent behaviors, by finding significant effects in some and non-significant effects in other studies. The question is whether some of these studies were underpowered.

Indeed, statistical power might be particularly hard to achieve in social networks studies that simultaneously consider network change (e.g., friendship relations) and change in individual attributes (e.g., the level of delinquency). The reason is the asymmetric number of observations on both levels: At each data wave, N nodes are connected through multiple network ties. When $\overline{k}$ is the average degree (it is typically larger than one in meaningful network studies) this results in $N \cdot \overline{k}$ observations. Nodal attributes are only observed N times per data wave (see Krivitsky and Kolaczyk, 2015). This implies generally lower information available in the estimation of behavior change mechanisms and in consequence also lower power to detect these mechanisms.

This paper introduces a procedure to estimate the statistical power of longitudinal

network studies that aim at employing SAOMs in the empirical analysis. It further aims at providing some guidelines for researchers who are designing new studies and raising awareness about critical issues such as missing data and participant turnover.

Typical decisions to be made regarding the *research design* of longitudinal social network studies include the size and delineation of a social network. In organizational studies, for example, researchers may decide whether to survey all members of the organization or only those affiliated to a specific department. A related decision is the number of social networks that are simultaneously studied. In studies situated in a school context, for example, researchers could face the decision of either collecting data in a single school or in multiple schools. Further research design decisions relate to the number and time intervals of data collection waves, the granularity of a behavioral variable, and potential restrictions of the number of peer nominations in network questions.

Besides design decisions, researchers may be uncertain about prevalence and magnitude of *social mechanisms* at play; those social mechanisms may partly be "interfering" in a sense that they reduce the power of the sample. Uncertainties could relate to the *type* of mechanisms (for example, are individuals influenced by their close friends or by all peers?) and to *effect sizes* (for example, by how much will the average risk to start smoking increase if an individual has a friend who smokes?). Besides detailing social mechanisms relating to the selection and influence processes, researchers should also consider other dynamic mechanisms that are present. For example, tendencies for reciprocity, clustering, and preferential attachment may affect the level of homophily in the network observed (Stadtfeld and Pentland, 2015). The same could happen through confounding processes, for example, homophily on a correlated variable. Finally, there may be interfering mechanisms that affect the power of the sample. Two of those are investigated in this paper: The first relates to processes generating non-response (e.g., due to non-consent / absence of participants), the second to processes generating turnover or dropout of participants between data collection waves.

The combination of possible research design decisions and varying assumptions about social mechanisms spans a potentially large space of *alternative scenarios*. The statistical power of those scenarios will vary significantly which implies that researchers should develop a good understanding as early as possible. The best time for the exploration of statistical power of alternative scenarios is thus in the planning phase of a longitudinal social network study. Acknowledging that an exhaustive coverage of this scenario space is hardly feasible, we propose a simulation-based procedure to assess power in longitudinal social network studies.

In section 2 we sketch a generic six-step procedure for social network power studies that is in line with classic power studies, for example, as introduced by Cohen (1977). Its relative complexity accounts for the fact that social network data distinguish themselves from simple random or clustered samples by a more complex structure. The procedure makes use of the R package NetSim (Stadtfeld, 2015) to simulate social network data, and of the R package RSiena (Ripley et al., 2016) to simulate and estimate SAOMs. Cohen (1977, p.4) describes that power in classical experimental studies depends on three parameters: the significance criterion, effect sizes and the sample size. Those parameters can also be found in our procedure in a somewhat extended fashion. The *significance criterion* will be set to a 95% confidence level for the model parameters of interest. For the *effect sizes* good a-priori are difficult, especially in view of the high interdependence between model parameters. Effect sizes relate to the degree to which social mechanisms are prevalent. For SAOMs, standardized effect sizes have not yet been developed, and therefore the parameters of the model must be used as effect sizes. Parameter interpretation is discussed by Snijders et al. (2010b, section 3.4). The *sample size* generalizes to multiple dimensions of th research design, such as the network size, the network delineation, and the number of data collection waves. The "sample size" is further affected by the prevalence of interfering mechanisms such as turnover rates and non-response.

To illustrate the six-step procedure, we discuss two fictitious research settings in

sections 3 and 4 that are inspired by what we perceive as "typical" empirical selection and influence studies. The first research setting in section 3 examines how the number of data collection waves and the boundaries of a network affect power. This research setting relates to exploring alternative *research design* decisions. The second research setting in section 4 discusses statistical power issues of a selection and influence study in a data collection involving multiple classrooms. In particular, we investigate to what extent statistical power is influenced by homophily and social influence effect sizes, by respondent data that are missing completely at random (Huisman and Steglich, 2008), and by turnover of students between data collection waves. This research setting relates to exploring a space of varying *social mechanisms* at play. Section 5 discusses the potential impact of this paper on the design of future longitudinal social network studies.

2 A procedure for the estimation of statistical power

The proposed procedure evaluates a range of alternative scenarios that vary in research designs and express uncertainty about the prevalence and magnitude of various (potentially interfering) mechanisms. The procedure is sketched in figure 1 and consists of six major steps.

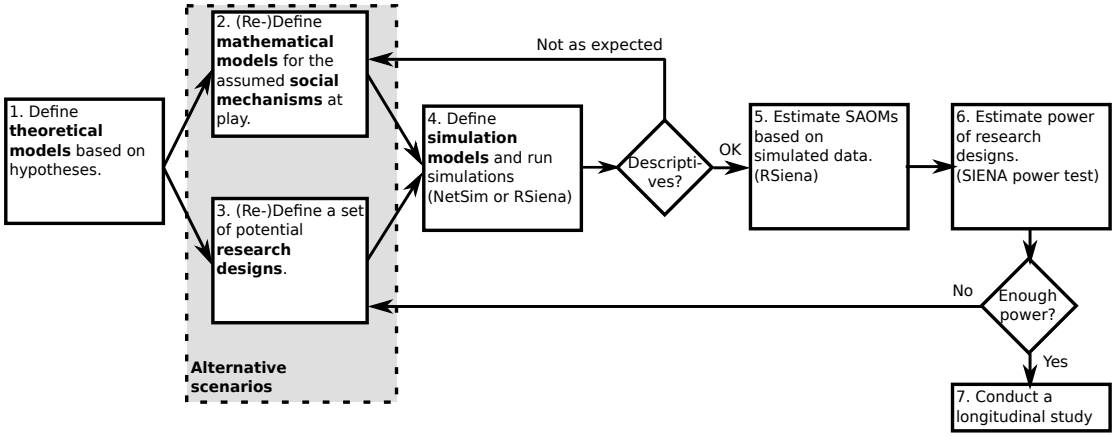
1. Define **theoretical models** based on hypotheses.
2. (Re-)Define **mathematical models** for the assumed **social mechanisms** at play.
3. (Re-)Define a set of potential **research designs**.
Alternative scenarios
4. Define **simulation models** and run simulations (NetSim or RSiena)
Not as expected
Descriptives?
OK
5. Estimate SAOMs based on simulated data. (RSiena)
6. Estimate power of research designs. (SIENA power test)
Enough power?
No
Yes
7. Conduct a longitudinal study

Figure 1: Overview of the procedure for the estimation of statistical power in longitudinal social network studies.

1. Each longitudinal social network study starts with the formulation of hypotheses on dynamic mechanisms. Typical hypotheses relate to homophily processes in the network formation (McPherson et al., 2001) and social influence processes on the attribute level (Friedkin, 1998). However, many other research questions in the domain of social networks can be considered that either relate to network change or to attribute change processes. The theoretical model is enriched with a number of other dynamic processes, such as reciprocity, transitivity, or popularity mechanisms (Kadushin, 2012).

→ The following two steps span the space of *alternative scenarios* that vary first, due to uncertainty about the *social mechanisms* at play (step 2) and, second, due to varying *research designs* (step 3).

2. The *social mechanisms* identified in step 1 are translated into formal *mathematical models*. The class of stochastic actor-oriented models (SAOMs) is a good starting point as it allows the combination of several network- and attribute-related social mechanisms (Snijders et al., 2010b; Steglich et al., 2010; Snijders and Steglich, 2015). But also other models could be applied as a mathematical framework, for example, tie-based Markov models that generate Exponential Random Graph distributions (Lusher et al., 2013, ch.12), or Hierarchical Latent Space Models (Sweet et al., 2003). It is possible that some aspects of the theoretical model cannot be expressed with SAOMs, for example, processes that lead to specific types of missing data or cause individuals to join and leave the population. Dynamic processes of that kind should also be formalized and added to the mathematical model. The research setting in section 4 focuses on the exploration of several mathematical models that express uncertainty about social mechanisms.

3. Potential *research designs* are defined to address the hypotheses formulated in step 1. A first ad-hoc attempt may build on designs of previous research studies.

Typical decisions in this step are the number of individuals in the study (i.e., number of networks or network boundaries), the number of waves of data collection, the time spans between subsequent waves, the granularity of a behavioral scale, or whether the number of nominations in a network questionnaire should be restricted. The research setting in section 3 focuses on the the problem of choosing between several feasible research designs.

4. Simulation models are defined for a reasonable subset of the alternative scenarios described by steps 2 and 3. Additional assumptions may be necessary. These may relate to starting distributions of individual attributes or network structures at the beginning of a data collection. For each simulation model a number of simulations is run (e.g., 200). Descriptives of the simulated networks and individual attributes should be checked at the end the simulations to avoid that the simulations generate unexpected or unrealistic outcomes. The "Goodness-of-fit" routine of the SIENA software (Ripley et al., 2016, sec. 5.11) can be used for this purpose. If descriptives of the simulated networks are unreasonable, the mathematical models from step 2 should be improved. In this paper, we simulate data with the R package NetSim (Stadtfeld, 2015) and the RSiena package (Ripley et al., 2016). Previous papers in which RSiena was applied in simulation studies are Snijders and Steglich (2015); Prell and Lo (2016). Example simulation scripts with RSiena and NetSim are published online[2].

5. The simulated data sets (say, 200 per simulation model) are used as data input for an estimation with the RSiena software. The models are specified according to the theoretical models in step 1. This step of re-estimating models may take a considerable amount of time as the number of simulation models is relatively large and the simulation-based estimation of parameters of the SIENA method is time-consuming.

6. For each simulation model, the percentage of cases is calculated in which sig-

nificant evidence supporting the hypotheses could be found in the re-estimation step 5. The significance can, for example, be tested at a $\alpha = 0.05$ significance level. A more efficient estimator could be given by estimating the mean and standard deviation of the parameter estimate or the mean of the t-ratios (with assumed variance 1) and estimate power from there. The percentage of positive findings is an estimate of the statistical power of the study design. If several research designs seem to have a good power, then the least costly can be chosen and the longitudinal study can be conducted. If the power in all research designs is too low, then extensions of the research design should be considered. This corresponds to updating the research designs in step 3.

3 Research setting 1: Opinion dynamics in four local communities

The first research setting discusses a (fictitious) research design in which the dynamics of friendship and opinion formation (negative – neutral – positive) in four local communities are observed. The communities are geo-spatially close to one another so that interpersonal ties may occur between them, however, ties within communities are more likely. We sketch a research study in which the social network and opinion dynamics of 120 individuals are of interest. The key hypotheses are that both homophily and influence processes with regards to opinions are prevalent. The design decisions take the network boundaries and the number of waves of data collection into account. To investigate the statistical power of different research designs, we follow the six-step procedure introduced in section 2.

3.1 Hypotheses and assumptions

In this study we are interested in two hypotheses, namely whether changes in opinions are explained by the opinions of friends (social influence) and whether individuals choose their friends based on opinion similarity (homophily). Several additional dynamic assumptions are made. First, individuals have a slight tendency for polarization. In disregard of social effects individuals are expected to converge to extreme opinions (negative or positive instead of neutral) in the long run. Second, we assume a social network formation that is partly explained by preferences for reciprocity, geo-spatial proximity (propinquity) and on the general preference for transitive structures. Third, personal networks of individuals are assumed to change faster than their opinions. Furthermore, we start with some straightforward assumptions about how the friendship network and the distribution of opinions look like at the beginning of the study.

3.2 Mathematical formulation

The hypotheses and the additional assumptions are formalized as a stochastic actor-oriented model (SAOM). Based on the parameters of "typical" empirical SIENA models (cp. Snijders et al. (2010b); Veenstra et al. (2013)), we formalize the exemplary model with the specification shown in table 1. The question how to translate hypotheses into SAOM parameters is nontrivial – empirical findings of similar studies (see, e.g. Snijders (2016)) can provide reasonable starting values. Therefore, it might be a good decision not to define just one model, but several models with varying parameters. This will reduce the risk that the scope of the power study is very limited by a choice of unreasonable parameters.

3.3 Research designs

We explore two scenarios. The first design decision is about the network delineation: Should data be collected in one, two or all four local communities (N = 30, 60 or 120)?

Mechanism	**SIENA effect name**	**Parameter**
Network change	`rate`	3.0
Density	`density`	-2.0
Reciprocity	`recip`	2.0
Transitivity	`transTrip`	0.2
Hierarchization	`cycle3`	-0.1
Propinquity (Distance)	`X`	-2.5
Homophily	`simX`	1.5
Attribute change	`rate`	0.6
Strive for polarization	`linear`	-0.8
	`quad`	0.2
Influence	`totSim`	0.8

Table 1: Specification of a stochastic actor-oriented model that expresses assumptions about the social mechanisms at play (step 2) in the first research setting. The focal mechanisms are emphasized.

The second design decision is concerned with the number of data collection waves. In this example, we consider collecting two waves, three waves or five waves of data. Obviously, the two decisions are directly related to the costs of the data collection.

3.4 Simulation models

We generate five simulation models based on the mathematical formulation and a subset of the space of potential research designs. The five simulation models relate to five research designs and are sketched in table 2. From each simulation model 200 data sets are simulated with the software package NetSim (Stadtfeld, 2015)[3].

Each simulation is based on an initial equal distribution of opinions and an initial friendship network. The starting network is simulated from an empty network with the stochastic actor-oriented model shown in table 1, except for the homophily and influence effects. After this initial process, individual attributes are randomly assigned to actors in order to achieve an initial observation in which network position and individ-

	one community (N = 120)	two communities (N = 60)	four communities (N = 30)
5 waves	X		X
3 waves		X	
2 waves	X		X

Table 2: Five out of nine possible simulation models are chosen.

ual attributes are uncorrelated – this relates to an assumption of the fictitious study that social effects on opinion formation only start playing out after the initial data collection. Figure 2 shows four networks that were extracted from one simulation run. Actors are positioned in a two-dimensional space; the distance between actors affects the propensity to form network ties. Checks of network densities and degree distributions reveal that the simulated networks are reasonable from a descriptive point of view. In particular, the simulation model is not degenerative in a sense that it would produce graphs with a density close to one in the long run. Therefore, we proceed with step 5 of the procedure. A visualization of a related dynamic four-community simulation can be found in the online appendix[4]. It demonstrates the non-degeneracy of the model specified.

3.5 Estimation with RSiena

After the simulations, the generated data are fitted to a stochastic actor-oriented model using the RSiena software. This model is specified with exactly the same parameters that were used in the mathematical model (see table 1). The simulation phase generated 1000 result sets (5×200) that consists of parameter estimates and standard errors. This process takes a significant amount of time (about one day on a standard personal computer) but can be accelerated by making use of parallel computing. All 1000 simulations and subsequent estimations with RSiena are independent and can be processed in parallel.

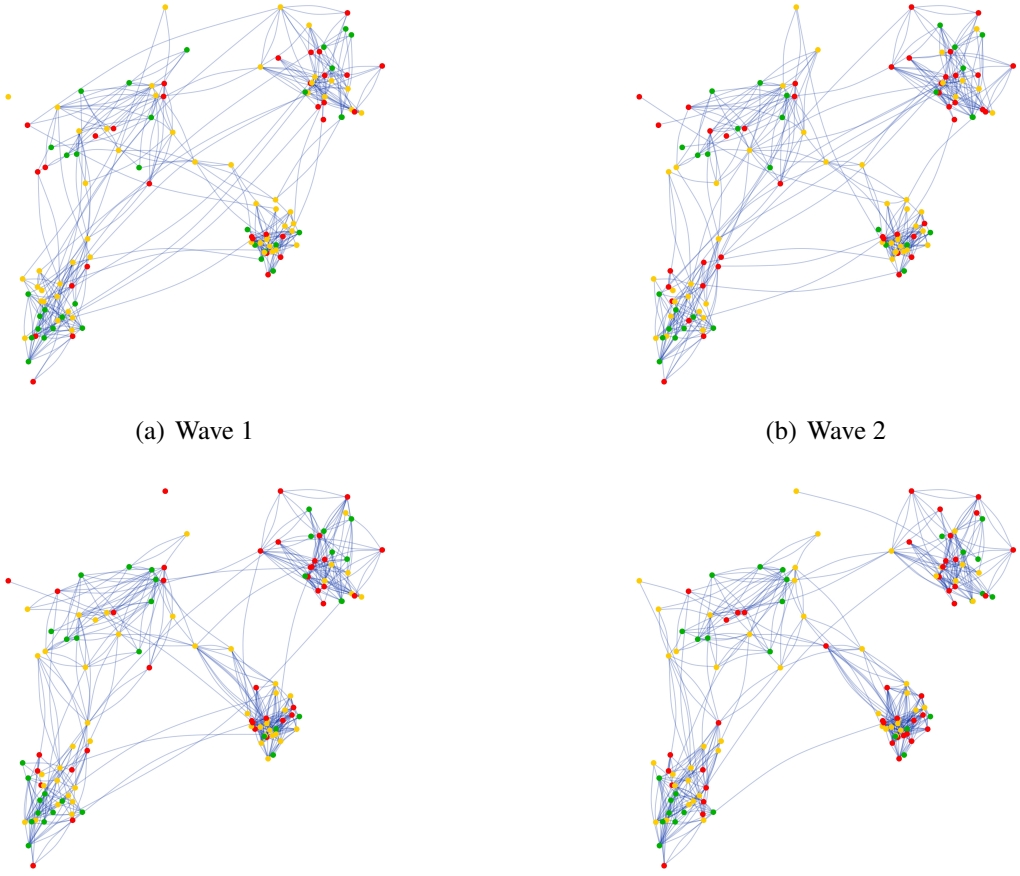

(a) Wave 1

(b) Wave 2

(c) Wave 3

(d) Wave 4

Figure 2: Four waves of data generated by the simulation process in one simulation run. Both the friendship network and the attributes (indicated by color codes) change over time following the mathematical model specified in table 1. All four local communities are shown. The network layout corresponds to the geo-spatial distribution of individuals in the study.

3.6 Evaluating the power

For each simulation model, the power of the parameters is evaluated. For example, the results of the scenario with two local communities (N = 60) and three waves of data collection are shown in table 3. It includes the effect names, the simulation model parameters (see table 1), the mean estimated parameters of 200 simulated data sets, their standard deviation and the power of the effects in this particular study design. The power column indicates the percentage of simulated data sets for which a parameter could be re-estimated significantly with a p-value smaller than 0.05.

Effects	**Sim. param.**	**Avg. est**	**St. dev**	**Power (%)**
Network change	3.0	2.49	0.32	
Density	-2.0	-3.17	0.19	100.0
Reciprocity	2.0	2.08	0.18	100.0
Transitivity	0.2	0.27	0.09	85.0
Hierarchization	-0.1	-0.19	0.16	17.5
Propinquity (Distance)	-2.5	-1.18	0.16	100.0
Homophily	1.5	1.55	0.40	99.5
Attribute change	0.6	0.58	0.21	
Strive for polarization (linear)	-0.8	-0.17	0.33	3.0
Strive for polarization (quad)	0.2	0.04	0.54	3.0
Influence	0.8	0.87	0.59	34.5

Table 3: Results of the power test for the simulation model with the data set reduced to N = 60 actors and 3 waves of data collection. The two parameters that relate to the hypotheses are highlighted gray.

The key parameters (homophily and influence) are highlighted gray. Homophily has a power of 99.5%, the influence effect a power of 34.5%. Assuming that the simulated mathematical models are indeed a good representation of the real social processes at play, we could only expect a statistically significant confirmation of the influence hypothesis in one out of three studies. This may not be a sufficiently good expectation. Note that some mean parameter estimates differ from the simulated values in table 3. This is due to the fact that the simulation model was specified and run on a complete network of 120 actors. Only after the simulation, a sub data set of 60 actors was extracted. This affects the estimates of the density, clustering- and distance-related parameters.

For example, propinquity matters less in this re-estimation that is based on just two communities. The density parameter, however, is more pronounced as it balances out the higher levels of network clustering and the smaller effect of the propinquity parameter. The parameter estimates are thus not unbiased in this example. Still, the power of most of these network-related effects is high. The power of the attribute shape effects (linear and quadratic) is very low which confirms that attribute-related effects are particularly prone to have a low statistical power.

A comparison of the power of the five research designs is given in table 4. The table now only focuses on the power estimates of the two key parameters homophily and influence that are related to the initial hypotheses. The columns express the research design decision about the network delineation which ranges from 120 actors (four communities) to 30 actors (one community). The rows show the varying number of data collection waves. The value in the table are again the percentages of models with significant results (in a 95% confidence interval) of the homophily (first value) and the social influence parameter (second value). These estimates of statistical power correspond to the last column in table 3.

		Community size					
		N = 120		**N = 60**		**N = 30**	
		Hom.	Inf.	Hom.	Inf.	Hom.	Inf.
Number of waves	**5 waves**	100	97.5			97.5	28.5
	3 waves			99.5	34.5		
	2 waves	99.5	34.5			34.5	10.0

Table 4: Percentage of significant findings (in a 95% confidence interval) of the homophily (first value) and the social influence parameter (second value) in five different cases in which sample size (number of local communities) and number of data collection waves vary. These power estimates are based on 200 simulations and re-estimations per parameter combination.

In the minimal design with two waves and 30 actors the expected power of the influence effect is only 10% and also the power of the homophily effect is rather low

(34.5%). The statistical power estimates of the three intermediate designs (120 actors and two waves, 60 actors and three waves, 30 actors and five waves) are similar to one another: The power of the homophily effect is very good and ranges between 97.5 and 99.5%, whereas the power of the influence effect is again low and ranges between 28.5 and 34.5%. Only the large research design with 120 actors and five waves of data collection has an excellent power of 100% for the homophily and 97.5% for the influence parameter.

3.7 Conclusions of the first power study

Based on the five research design evaluations, researchers could now decide on how to conduct the longitudinal study on opinion and network formation in the four local communities. The small-scale research design seems not to be adequate in any case. If the influence hypothesis was of less interest, the most feasible of the three intermediate research designs could be chosen. Only the large research design promises good statistical power for the estimation of both homophily and influence effects. To elaborate on the power of the influence effect, researchers might want to run further power studies with, for example, 120 actors and three waves, 60 actors and four waves, or 90 actors and three waves. This would mean going back to step 3 (define a set of potential research designs) of the six-step procedure.

4 Research setting 2: Co-evolution of friendship and delinquency in 21 schools

In the second research setting, we investigate how varying effect sizes, missing data and change in the composition of study participants may affect the power of selection and influence effects. We choose a setting that resembles a typical longitudinal network study in a population of schools and is inspired by the study of Baerveldt et al. (2008)

on friendship selection and delinquency. We conduct a power study based on initially observed networks and delinquency attributes in wave 1 and simulate one additional wave of data based on an empirically estimated model that is similar to the on in the original study. Compared to the first research setting, the number of participants is very high (N = 742 students, distributed over 21 schools). The simulated data are based on three empirically observed school networks and delinquency scores of 106 school students that are replicated seven times each in order to construct an artificial sample with 21 schools and 742 students. Schools have 33, 36, and 37 students. This study focuses on how effect sizes (absolute values of simulated effects), participant turnover (students leaving and students joining the population between waves) and missing data (students not answering the questionnaire) affect the statistical power of the study design. We again follow the six-step procedure proposed in section 2.

4.1 Hypotheses and assumptions

The key hypotheses are that both homophily and social influence processes regarding delinquency are prevalent within school networks. In particular, we are interested in the effect of individuals selecting friends who are similar regarding the level of delinquency (homophily) as well as influence effects regarding delinquency. As in research setting 1, we further assume the presence of a number of dynamic network mechanisms (e.g., reciprocity, transitivity, gender homophily). Besides those we now expect processes that result in student turnover between data collection waves and missing data through non-participation. The rate of missing data and student turnover is assumed to be uncertain in the design phase of the study.

4.2 Mathematical formulation

We use the stochastic actor-oriented model to describe changes in the network structure and the individual delinquency variables. In the mathematical formulation we follow the

empirical model of Baerveldt et al. (2008, table 5, p.574) with some adaptations. Two effects regarding general importance of school friends and ethnical homophily are omitted for reasons of simplicity. An effect capturing an interaction between reciprocity and transitivity (Reciprocity in triads, see Block, 2015) is added to the friendship model and a quadratic shape effect is included in the behavior change part of the model. Thereby, the model is closer to state-of-the-art SAOM specifications[5]. The complete specification of the SAOM is shown in table 5. The parameters used for the simulation model were estimated on an empirical sample of 10 school classes using a meta-analysis (Snijders and Baerveldt, 2003). The focal parameters are highlighted gray. We test the power of two models: One in which we simulate the empirically estimated effect parameters ("low" parameters), and one in which we use slightly higher parameters.

Mechanism	**SIENA effect name**	**Parameter**
Network change	`rate`	4.3
Density	`density`	-3.1
Reciprocity	`recip`	2.4
Transitivity	`transTrip`	1.2
Reciprocity in triads	`transRecTrip`	-0.8
Homophily Sex	`sameX`	0.6
Homophily Delinquency	`simX`	low: 0.4 / high: 0.6
Delinquency change rate	`rate`	1.3
Strive for delinquency	`linear`	-0.2
	`quad`	-0.2
Influence Delinquency	`avAlt`	low: 0.3 / high: 0.4

Table 5: Formal specification of the mathematical model in the second research settings. The focal mechanisms are highlighted gray.

This basic model is extended by two straightforward mechanisms. The first mechanism describes turnover of students after half of the data collection period, the second mechanism generates missing data that stems from completely random non-participation of some students in the two data waves (one empirical, one simulated).

The turnover mechanism explains how students leave and join the sample. At half-time between the two data collection waves, a fixed number of students drops out of each school (0, 1, or 3). At the same time, the same number of students joins the school so that the school size (ranging from 33 to 37 individuals) remains constant. The new students are network isolates in the moment they join the school and only then start forming friendship relations. The attributes of a new student are randomly chosen based on a cross-frequency table of the attributes of all students (gender x delinquency) in the population at the time when the student turnover occurs.

The missing data mechanism relates to random non-participation in a survey wave. From both data collection wave a fixed number of students is chosen in each of the seven schools (0, 1, 3, 5, 7) of whom network nominations and delinquency levels are deleted (missing data). The number of missing entries is the same in both data collection waves. The two random draws of missing individuals in the two waves are independent.

In this research setting we thus assume uncertainty about the levels of student turnover (0, 1, 3 $\widehat{=}$ 0%, 2.8%, 8.5%), missing data (0, 1, 3, 5, 7 $\widehat{=}$ 0%, 2.8%, 8.5%, 14.2%, 19.8%), and the strength of homophily (simX in $\{0.4, 0.6\}$) and influence mechanisms (avAlt in $\{0.3, 0.4\}$).

4.3 Potential research designs

We do not consider different research designs. The statistical power of the study is tested for a research design that includes all 21 schools (N = 742 students), two waves of data collection, binary friendship nominations and a five-point delinquency scale. The space of alternative scenarios is therefore only spanned by the rates of missing data, student turnover rates, and the strength of selection and influence mechanisms.

4.4 Simulation models

The simulation models are based on the parameters in table 5 (one model with low and one with high homophily and influence parameters) and all 15 combinations of student turnover rates and missing rates. Each simulation model is simulated 200 times with the RSiena software (Ripley et al., 2016). An R function was developed for the simulations conducted in this study that combines RSiena-based simulations with the interfering processes of participant turnover and missing data. The first wave of data is taken from the empirical data of Baerveldt et al. (2008). A second data wave is simulated for each school class separately. In total, 126,000 data sets are generated.

The data have certain particularities that are illustrated in table 6. The average degree is very low (1.4 ties, the maximum in-degree is 5) even though the school networks are relatively big (33, 36, and 37 individuals). The average level of delinquency is 1.8 on a scale that ranges from 0 to 4. The dispersion of delinquency values is low. Of 742 individuals only 56 (7.5%) have a minimum score of 0, and 21 (2.1%) have a maximum score of 4.

	mean	**min**	**max**
friendship in-degree	1.4	0	5
delinquency	1.8	0	4

Table 6: Descriptives of the empirical data in research setting 2

After conducting simulations, we check the goodness of fit (Ripley et al., 2016) of a sub sample of the simulated networks regarding degree distributions and triad census and compared those to the empirically observed second data wave. The simulated networks are found to be similar to the empirical networks by which we conclude that the simulation model are appropriate.

4.5 Estimation with RSiena

The 126,000 simulated data sets are grouped into sets of 21 that constitute our empirical sample size. Then, parameters are estimated with the RSiena software using the

"multigroup" option (Ripley et al., 2016, section 11.1) for the analysis of multiple networks. The re-estimation of sets of 200 simulated multigroup data sets (that all share the same model specification, turnover rates, and missing rates) takes between one and eight hours on a computer with 24 cpus. A computer cluster has been used for this step so that multiple SIENA re-estimations could be run in parallel.

4.6 Evaluating the power

The power estimates are given in figures 3 and 4. Figure 3 shows the power estimates for the homophily and the influence parameter of the "weak" model in table 5, figure 4 those of the model with stronger effects. Three lines indicate power of turnover rates of 0%, 2.8% and 8.5%. The x-axis covers different missing rates. A dotted line at the 0.05 level indicates the chosen significance level that would be the expected power of unbiased estimates that have no information value at all (zero effects). In both models, the power rates with no turnover and 2.8% (low) turnover develop somewhat similar. There is a significant drop in power with turnover rates of 8.5%. On problem that we encounter is that it is more difficult to achieve convergence of the estimation routine (Ripley et al., 2016, sec.6) in case of models with an 8.5% turnover rate. While close to 100% of the models with zero and 2.8% turnover converged, convergence could only be achieved in about 80% of the high-turnover models. The coverage rates under the null are almost all sufficiently close to 0.95 to conclude that under the null hypothesis the distribution of the parameter estimates is very close to a normal distribution with mean 0 and standard deviations equal to the reported standard error. The exception is the estimated social influence parameter (avAlt) in case of high-turnover (8.5%) models, where the standard errors are inflated. With the small remaining sample size and the skewed dependent variable, this may be due to the occurrence of the so-called Donner-Hauck phenomenon (Hauck and Donner, 1977; Ripley et al., 2016, sec.8.1) where the standard error is inflated and the Wald test should not be used for hypothesis testing.

The very low rejection rates under the null are associated with lower power for the Wald test, if it would be used. This explains why the power of the high-turnover models drops below the 5% line in figures 3 and 4. From a design point of view, the interpretation of the results is clear: with this amount of turnover for only two waves of data, it is impossible to have a satisfactory study of social influence. In the following, we discuss results of the models in which the turnover rate was zero or 2.8%.

In the models with weaker effects (figure 3), the power of the homophily parameter and the influence parameter are rather low. The maximum power in a model without turnover and missings is 30% (homophily, simX) and 38% (social influence, avAlt). When the missing rates increase to 19.8% the power of the homophily parameter drops to the random expectation of a null effect when a p-value of 5% is chosen (5% power). The power of the influence effect remains only slightly higher.

The models with more pronounced effects (figure 4) start off from higher power values. In case of no missing and no turnover the power of the homophily effect is 53%, the power of the influence effect is 70%. A turnover rate of 2.8% seems not to affect the power estimates a lot. The lines related to no and little turnover in the figures are again somewhat similar. In a model with 19.8% missing rates, the statistical power drops to 19% and 22% for homophily and social influence respectively.

4.7 Conclusions of the second power study

The second case study illustrates the effect of turnover and missing rates on the power of a longitudinal study design. Missing data of 19.8% (the highest simulated value) reduces the power significantly. The power of the influence parameter in the model with weaker parameters, for example, dropped from 37.5% to 7.5%. The latter is close to the expectation of 5% for estimates with no information value. Turnover also has a negative effect on power and we further observed an inflation of standard errors, probably due to the so-called Donner-Hauck phenomen. It turned out that with just two waves of

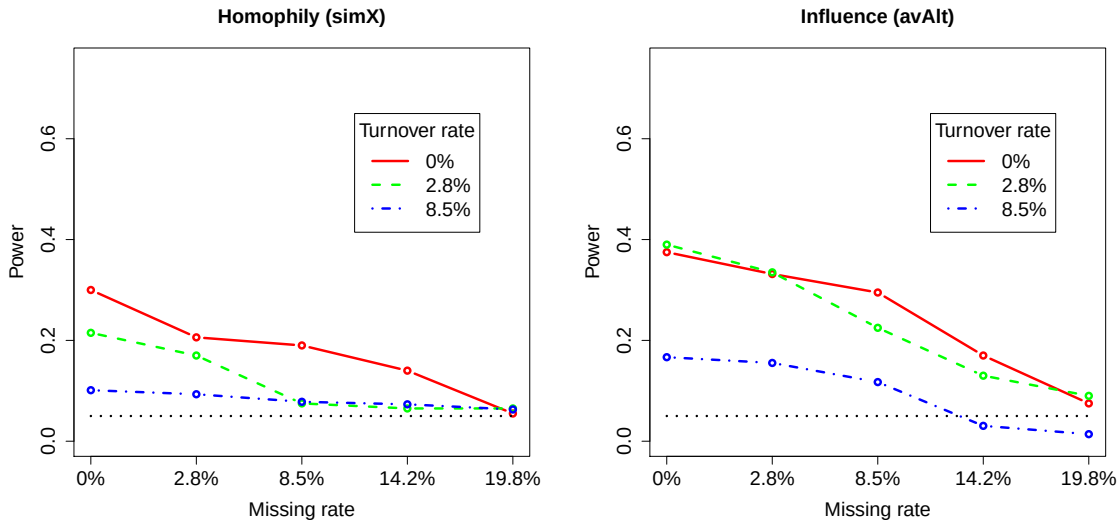

Figure 3: Power of models with the weaker homophily (simX = 0.4) and influence (avAlt = 0.3) parameters. Missing rates are indicated in the x axis, turnover rates are given by the three lines. The black dotted line indicates the chosen significance level that would be the expected power of unbiased estimates that have no information value at all (zero effects).

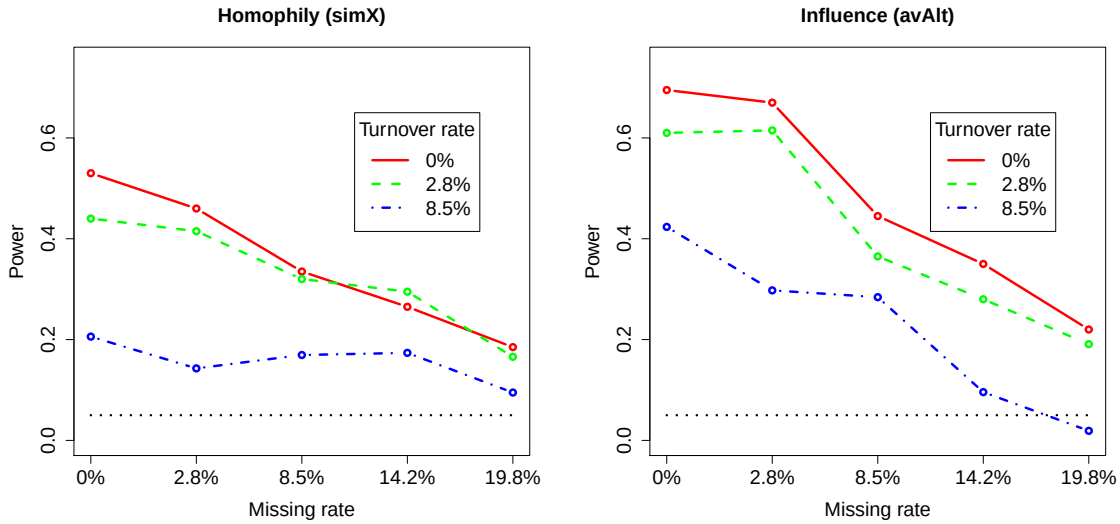

Figure 4: Power of models with with the stronger homophily (simX = 0.6) and influence (avAlt = 0.4) parameters. Missing rates are indicated in the x axis, turnover rates are given by the three lines. The black dotted line indicates the chosen significance level that would be the expected power of unbiased estimates that have no information value at all (zero effects).

data and a turnover ratio of 8.5% the statistical power was unsatisfying in almost all simulation models.

A notable observation is further that the power of the homophily parameter is generally lower than the power of the influence parameter. This seems counterintuitive given our initial discussion that homophily inference is based on $N \cdot \bar{k}$ observations while influence effects are estimated based on N observations per wave. In this example, however, we use data with specific particularities that probably strongly affect the power of the study. First, the network is very sparse. Initially, only 1.4 friendship nominations exist ($\bar{k} = 1.4$). At the same time, we estimate a higher number of effects in the network change sub-model (seven as compared to four in the behavior change sub-model). Second, the dispersion of the delinquency variable is very low; only 7.5% and 2.1% of individuals were in the lowest and highest category of the five-point scale in the first data wave. The homophily and the influence parameter are estimated based on cross-lagged statistics (Steglich et al., 2010) that do not carry a lot of information when the variable dispersion is low and only few ties are observed. Researchers facing this problem might for example want to consider using a more fine-grained delinquency scale that generates a higher dispersion. This might significantly improve the power of the homophily parameter in particular. As an improved estimation strategy it should be considered to use a maximum likelihood routine (Snijders et al., 2010a) as it uses information more efficiently and should thus provide better power expectations. Using maximum likelihood estimations in the re-estimation of simulated models (step five of the six-step routine) is possible in general but will take significantly more time.

5 Discussion and conclusions

In this paper, we presented a novel procedure for the estimation of statistical power in longitudinal social network studies. In particular, we discussed study designs that aim at investigating dynamic selection and influence mechanisms with stochastic actor-

oriented models (SAOMs). An increasing number of empirical studies of that type have been published recently (see Veenstra et al. (2013) for a partial overview). Those studies report mixed findings about homophily and social influence processes which we argued might be related to statistical power issues. So far, however, not a lot is known about appropriate research designs and statistical power. The six-step procedure that we presented in this paper can be seen as a tool for the investigation and comparison of statistical power of longitudinal social network research designs. We demonstrated its utility in two extensive research settings that focused on the effect of network size, number of data collection waves, effect sizes, missing data, and participant turnover on statistical power.

The two research settings that we presented did not aim at providing practical rules of thumb in the first place. However, it became clear that network delineation, number of data collection waves, effect sizes, missing data and participant turnover may strongly affect the power of longitudinal selection and influence studies. In research setting 1 (section 3), we specified a mathematical model of selection and influence with pronounced effect sizes. Yet, a simulated small-scale research design with just 30 individuals and two waves of data collection was found to be not appropriate for empirically testing either of the two effects. A research design with five waves of data and 120 individuals provided excellent power expectations for both the homophily and the influence effect. In research setting 2 (section 4), we specified a similar mathematical model for selection and influence dynamics among 742 students distributed over 21 schools. The simulated effect sizes in this study were smaller, we only simulated two data waves, and the initial data carried a lot less information. Given those specificities, we found that a missing data rate of 19.8% would significantly reduced the power of homophily and influence parameters. In a simulation model with weak effect sizes, the power was practically indistinguishable from the power of parameters that carry no information at all (null effects). A turnover rate of 8.5% also had a strongly negative effect on power. A practical issue that arose in models with high participant turnover is that it is harder

achieve convergence in the estimation routine. Missing data and student turnover rates in that magnitude are not uncommon. This underlines the importance of social network data collections that aim at high participation rates and panel stability over time.

The two empirical settings provide some intuition about issues that researchers should be concerned about, however, the quantitative results should not be generalized. We could indeed show that in these cases the power estimates are highly affected by variations in a number of research design and social mechanism parameters. Those parameters jointly affect the power. For example, the level of participant turnover where it has serious effects on power will depend on the distribution of variables and the network structure.

A common question that we were confronted with when presenting this work is whether the procedure may be used to investigate if insignificant effects in an empirical study result from a lack in statistical power. Indeed, it is common sense among statisticians that post-hoc power studies are *irrelevant in the interpretation of empirical results* (Cox, 1958; Goodman and Berlin, 1994; Senn, 2002; Lenth, 2007). Estimating the power of a research design as a result of not finding significant evidence for a hypothesis may lead to the dangerous conclusion that evidence for a (non-significant) social mechanism may just not have been found because of a lack of power. Yet, the level of confidence about an estimate is already captured by the estimated standard errors or confidence intervals.

Post-hoc power studies should thus not be used in the interpretation of parameters. However, they may motivate future research in case they suggest that certain adaptations may indeed improve the power of a study design. Gelman and Carlin (2013) propose that post-hoc "design analyses" may generally be useful when assumptions about social mechanisms at play stem from prior expectations or prior empirical findings but not from the empirical estimates.

The six-step procedure proposed can provide new guidelines for the design of longitudinal social network studies. We hope that it inspires systematic investigation of

longitudinal research designs on a number of dimensions. As examples, we could show that network size, duration of a study, effect sizes, missing data and participant turnover matter for the statistical power of a study. Other directions are to be explored in the future: How do, for example, measurement scales, systematic types of missing data, varying assumptions about interfering social mechanisms, alternative influence mechanisms, measurement errors, and varying time intervals affect the power of a study design? Many of these topics are of critical importance for empirical research and should thus be explored in varying contexts in the future. The six-step procedure that we presented in this article provides an important toolbox to do this.

ENDNOTES

1 A nearly complete list of SAOM-related publications is available at Snijders (2016)

2 Scripts and supplementary material will be published online with publication of the paper.

3 In this example in which the mathematical model is completely describable within the SAOM framework, the RSiena software could have been used for simulations as well

4 A simulation video is published with this paper.

5 The model of Baerveldt et al. (2008) is flawed because it omits the quadratic shape parameter that models dispersion of the behavioral variable. What they find as influence is essentially underdispersion that was not captured and hence appears as "staying close to friends" for a lack of closer effect in the model.

REFERENCES

Baerveldt, Chris, Beate Völker, and Ronan van Rossem. 2008. "Revisiting Selection and Influence: An Inquiry into the Friendship Networks of High School Students and Their Association with Delinquency." *Canadian Journal of Criminology and Criminal Justice* 50:559–587.

Block, Per. 2015. "Reciprocity, transitivity, and the mysterious three-cycle." *Social Networks* 40:163–173.

Cohen, Jacob. 1977. *Statistical Power Analysis for the Behavioral Sciences*. New York: Academic Press Inc. first published in 1969.

Cox, David Roxbee. 1958. *Planning of experiments*. New York: John Wiley.

Friedkin, Noah E. 1998. *A Structural Theory of Social Influence*. Cambridge University Press.

Gelman, Andrew and John Carlin. 2013. "Retrospective design analysis using external information." `http://www.stat.columbia.edu/~gelman/research/unpublished/retropower5.pdf`; last visited: June 2014.

Goodman, Steven N. and Jesse A. Berlin. 1994. "The Use of Predicted Confidence Intervals When Planning Experiments and the Misuse of Power When Interpreting Results." *Annals of Internal Medicine* 121:200–206.

Hauck, W.W.J. and A. Donner. 1977. "Wald's test as applied to hypotheses in logit analysis." *Journal of the American Statistical Association* 72:851–853.

Huisman, M. and C. Steglich. 2008. "Treatment of non-response in longitudinal network data." *Social Networks* 30:297–308.

Kadushin, Charles. 2012. *Understanding Social Networks. Theories, Concepts and Findings*. New York: Oxford University Press.

Krivitsky, Pavel N. and Eric D. Kolaczyk. 2015. "On the Question of Effective Sample Size in Network Modeling: An Asymptotic Inquiry." *Statistical Science* 30:184–198.

Lenth, Russell V. 2007. "Statistical power calculations." *Journal of Animal Science* 85:E24–E29.

Lusher, Dean, Johan Koskinen, and Garry Robins (eds.). 2013. *Exponential Random Graph Models for Social Networks*. Cambrisge University Press.

McPherson, J. Miller, Lynn Smith-Lovin, and John Manuel Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Annual Review of Sociology* 27:415–444.

Prell, Christina and Yi-Jung Lo. 2016. "Network formation and knowledge gains." *The Journal of Mathematical Sociology* 40.

Ripley, Ruth M., Tom A. B. Snijders, Zsófia Boda, András Vörös, and Paulina Preciado Lopez. 2016. *Manual for SIENA 4.0.* Nuffield College and Department of Statistics, University of Oxford. Version: December 6, 2013.

Senn, Stephen J. 2002. "Power is indeed irrelevant in interpreting completed studies." *British Medical Journal* 325:1304–1304. Letter.

Snijders, Tom A. B. 2016. "The Siena webpage." `http://www.stats.ox.ac.uk/~snijders/siena/`.

Snijders, T. A. B. and Chris Baerveldt. 2003. "A Multilevel Network Study of the Effects of Delinquent Behavior on Friendship Evolution." *Journal of Mathematical Sociology* 27:123–151.

Snijders, Tom A. B., Johan Koskinen, and Michael Schweinberger. 2010a. "Maximum likelihood estimation for social network dynamics." *The Annals of Applied Statistics* 4:567–588.

Snijders, T. A. B. and C. Steglich. 2015. "Representing Micro–Macro Linkages by Actor-based Dynamic Network Models." *Sociological Methods & Research* 44:222–271.

Snijders, Tom A. B., Gerhard G. van de Bunt, and Christian E. G. Steglich. 2010b. "Introduction to Stochastic Actor-based Models for Network Dynamics." *Social Networks* 32:44– 60. Dynamics of Social Networks.

Stadtfeld, Christoph. 2015. "NetSim: A Social Networks Simulation Tool in R." *R package vignette* `http://www.social-networks.ethz.ch/research/research-projects.html`.

Stadtfeld, Christoph and A. Pentland. 2015. "Partnership Ties Shape Friendship Networks: A Dynamic Social Network Study." *Social Forces* 94:453–477.

Steglich, Christian, Tom A. B. Snijders, and Michael Pearson. 2010. "Dynamic Networks and Behavior: Separating Selection from Influence." *Sociological Methodology* 40:329–393.

Sweet, Tracy M., Andrew C. Thomas, and Brian W. Junker. 2003. "Hierarchical Network Models for Education Research: Hierarchical Latent Space Models." *Journal of Educational and Behavioral Statistics* 38:295–318.

Veenstra, René, Jan Kornelis Dijkstra, Christian Steglich, and Maarten H. W. van Zalk. 2013. "Network-Behavior Dynamics." *Journal of Research on Adolescence* 23:399–412.